EBOOK

Modern Data Protection for Databases

rubrik

Table of Contents

Introduction

The pace of digital transformation will accelerate further in 2021, as more and more economic activity moves online. The increased speed of business magnifies the pressure on busy DBAs and operations teams, who are already working hard to address increasing demands without sacrificing security or data protection. Your team may find it hard to maintain existing backup service levels as data volumes increase and databases sprawl across on-premises and cloud—especially given the growing demands from secondary users like developers that want fast access to data. Tensions between backup admins and DBAs are also increasing due to conflict around data ownership, leading to increased costs and declining productivity.

Many organizations are addressing these problems by modernizing their approach to data protection. Modern database protection efficiently bridges the divide between DBAs and backup administrators, making it simpler to provide the right level of protection for every database, and ensures that nothing falls through the cracks as your database environment becomes more dynamic.

A modern data protection solution allows you to:

- **Reduce manual work.** Your team already has enough on its plate. Modern database protection solutions automate routine tasks, easing the burden on staff *and* reducing the risk of errors that leave data unprotected.
- **Control recoveries.** Most DBAs are happy to give up responsibility for regular backups but want to maintain control over recoveries. A modern data protection solution allows you to delegate responsibility for routine data protection tasks to backup administrators while giving DBAs the visibility and control to accomplish recoveries quickly.
- **Ensure recoverability.** There's no worse feeling than initiating an important recovery and having it fail. Modern data protection ensures that backups complete successfully and remain recoverable when you need them.
- **Gain access to the cloud.** With data volumes and compliance regulations increasing, most organizations are looking at not only protecting data natively in the cloud but leveraging the cloud for cost-effective long-term archival. Modern data protection integrates seamlessly with the cloud, streamlining backup and archival across environments.

This guide explores what modern database protection entails, examining the difference between the legacy and modern approach in four areas essential for successful data protection.

Advantages of Modern Database Protection

Because databases power mission-critical applications, backups must be easily accessible and quickly recoverable. With traditional data protection approaches, unfinished backups can go undetected, database refreshes may take days, and manual scripting is necessary to enable multi-step workflows. When dealing with large, highly-transactional databases, you run the risk of lengthy RTOs and data loss. A full restore may be necessary just to query a few tables.

Modern data protection solutions address these challenges, simplifying data protection for databases while offering near-zero RTOs and rapid database clones to support application development and other needs. Policy-driven methods eliminate painful scripting and job scheduling, enabling your team to manage large-scale, dynamic environments with automated database discovery and protection that satisfies your established SLAs—dynamically protecting databases as they are created.

A modern solution enables you to manage all your physical and virtualized databases from one interface—both on-premises and in the cloud—and easily integrate with your existing workflows or management tools. The key elements of a modern data protection solution include:

- **Automation.** That eliminates manual work, reduces human error, and—more importantly—gives you time back.
- **Near-zero RTOs.** Ensuring you can meet or exceed stringent SLAs regardless of the size of the database.
- **Advanced security.** To protect backup data against increasing threats including ransomware.
- **API integration.** Enabling data protection to integrate more easily with existing methods and tools.

The sections that follow compare and contrast traditional methods for protecting databases with modern data protection, demonstrating how your operations can benefit.

Managing Discovery and Data Protection

With the number of database instances growing and environments becoming more dynamic, it's becoming harder to keep up with growing data protection tasks.

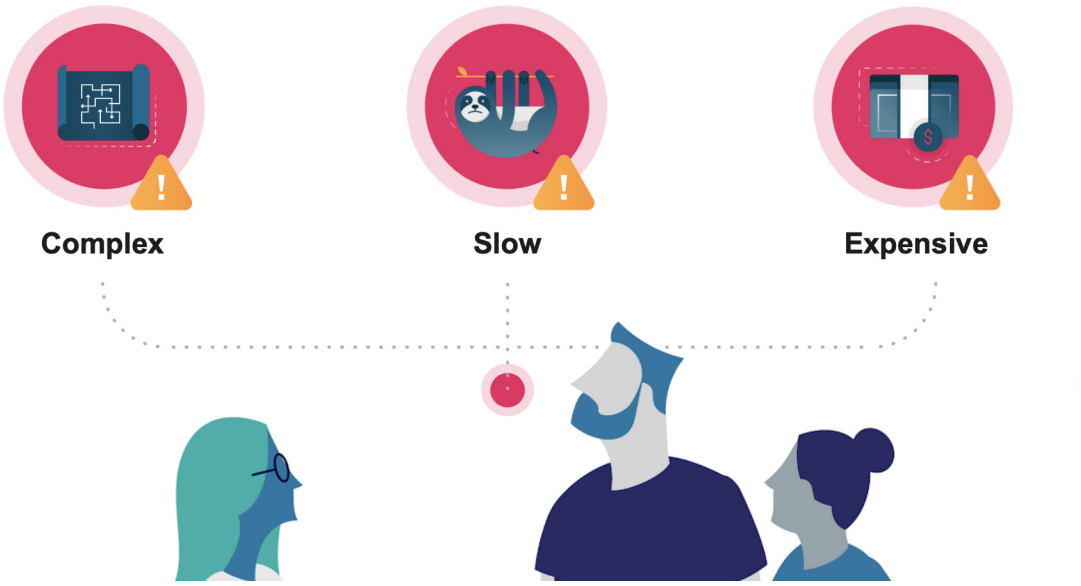

Legacy

With legacy data protection, adding a new database to your backup schedule is a multi-step, manual process. Each new backup job has to be created using native scripting, which not only requires writing the script (or modifying an existing one), but also installing it on the server and scheduling it to run—for instance using *cron* on Linux or the SQL Server Agent job scheduler. You also have to identify a storage target with enough capacity to satisfy your retention schedule; you may have to negotiate with a storage admin to get the necessary space, adding more friction to the task.

None of this is that difficult with one or a few database instances, but when you're operating at scale, the challenges mount quickly. For instance, backup scheduling has to be coordinated to ensure that you don't end up with too many backups running at the same time, creating bottlenecks on servers, storage, or networks. Simple scheduling mistakes can result in:

- Backups taking too long, potentially slowing down production
- Backups failing, putting data—and future recoveries—at risk
- Extended troubleshooting to identify and correct errors

When it comes to new databases, DBAs have to keep up with what's happening on database servers at all times to ensure that databases receive appropriate protection. While it's possible to write dynamic scripts to detect new instances and send an alert, it's difficult to guarantee that nothing slips through.

Modern

The right modern data protection solution automates the manual tasks associated with database discovery and protection. Once the software knows which servers run databases, it automatically discovers new instances, determines configurations, and protects them in a consistent manner based on policy, eliminating the need for error-prone and time-consuming manual scripting and scheduling.

An optimized backup schedule is created and executed automatically—according to your operational constraints—eliminating the risk of too many jobs running at once and causing problems. Modern data protection encompasses backup, replication, and archival, with the ability to utilize the cloud for long-term retention—all in a single policy.

Managing Recoveries

Because critical applications are built and hosted on databases, downtime has a direct impact on a company's revenue. Time is of the essence when it comes to database recoveries. You have to ensure that critical databases will be back online as quickly as possible should an outage occur.

Legacy

With traditional recovery, there are two fundamental challenges:

- You have to have a place for data to land
- It takes time to physically copy data

The database restore process is closely tied to the backup process. Whatever your backup process, it has to be unwound to execute a restore. You first have to determine whether to restore the most recent backup or an older one (that you have to then locate). You need a script to perform the restore, and you also have to have access to the necessary file(s).

If the necessary files are not online—perhaps you retain older backups on tape—a backup admin may need to restore from tape before a DBA can begin the recovery. Because there are a lot of variables, recoveries often take longer than expected or desired. Once a recovery does begin, a lot of data may have to be physically copied before the backup logs can be replayed to bring the database up to date and back online.

Given the critical nature of database applications, DBAs leverage other tools ahead of backup to ensure high availability. However, backups are often the last line of defense; you never want to find out that you can't restore a database backup in the middle of a crisis. You need methods to ensure that backups are completed successfully, and you also have to test your restore processes periodically to ensure that everything continues to work as expected.

Modern

The right modern data protection solution not only automates backup processes, it provides more flexible recovery options, enabling DBAs to devote their time to other tasks without relinquishing control over restores, especially in critical situations. Modern data protection simplifies critical database recoveries as well as the more common recovery scenarios that DBAs perform on a daily basis. Role-based access control (RBAC) gives the backup team supervision over backups while ensuring that DBAs maintain full control over critical recoveries. Multiple automated options for recovery address diverse use cases.

Simplified Recovery

One of the most stressful tasks for DBAs is recovering a production database in place, overwriting database files on the production host. This may only be necessary when there has been a complete database failure. Modern data protection automates this process, minimizing the stress on DBAs. A complete database can be restored directly to the original host from online backups and recovered to the latest consistent state, or the recovery can be customized to restore the database to any desired point in time.

Near-Zero RTO

For even faster recovery, a modern data protection solution allows you to quickly clone the backup of the production database and mount it directly from backup storage to the production host, eliminating the need to copy files. This approach reduces the recovery time to near zero, irrespective of the size of the database. Once the database is back in production, files can be migrated to the production host in the background to restore normal operation.

Flexible Options

Modern data protection provides a flexible recovery toolkit, making it quick and easy to address any recovery scenario and meet ad hoc requests efficiently:

- Copy database files directly to the original host in the event of a production recovery
- Easily copy your production database onto an alternate host
- Mount any backup to an alternate host for immediate data access or to pull out the data you need
- Mount any backup to the original host for instant recovery during a production outage
- Access a files-only recovery to manage restores using pre-existing scripts

Servicing the Data Needs of Users

As digital transformation accelerates, companies focus on greater business insights and new digital services, increasing the demands on DBAs. The tasks that DBAs perform on a regular basis to serve secondary data user needs—creating a reporting copy, pushing a copy for QA or development work, and so on—revolve around backup and recovery processes. The ease or difficulty of performing these tasks has a direct impact on DBA and user productivity—and the bottom line.

Legacy

As you saw in the previous section, with traditional data protection methods restores can be time-consuming. To service a single user request for a database copy, a DBA has to identify the right version of the file(s) to restore, create a script to do the restore (taking into account any special user requirements), and execute the script to carry out the process. Because a restore results in a full data copy, a DBA may also spend significant time making sure enough storage space is available—and negotiating with a storage admin if it's not. Storage often becomes the limiting factor, determining how many database copies you can maintain and limiting the speed of your business.

And once a secondary user has finished a project, someone has to clean up and recover the storage. If you bought extra storage in order to support a special project, you could end up with expensive equipment sitting idle when it's over.

Modern

Choosing a modern data protection solution enables you to unlock the power of your backups through automation. Instead of having to assemble scripts or build out detailed processes, you simply pick a point in time, and the software executes the restore, creating the necessary database copy quickly and reliably.

Empower Secondary Users with Self Service

With traditional methods, users lack the knowledge (and permissions) to execute a restore. They have to submit a request and wait for a DBA to carry it out. With modern data protection, a business analyst, developer, or other user can take advantage of self service capabilities to create the database copies they need. Full APIs make it easy to automate tasks like cloning and refreshing database copies, allowing users to select these options from a service catalog and get what they need immediately—instead of filing a ticket and crossing their fingers.

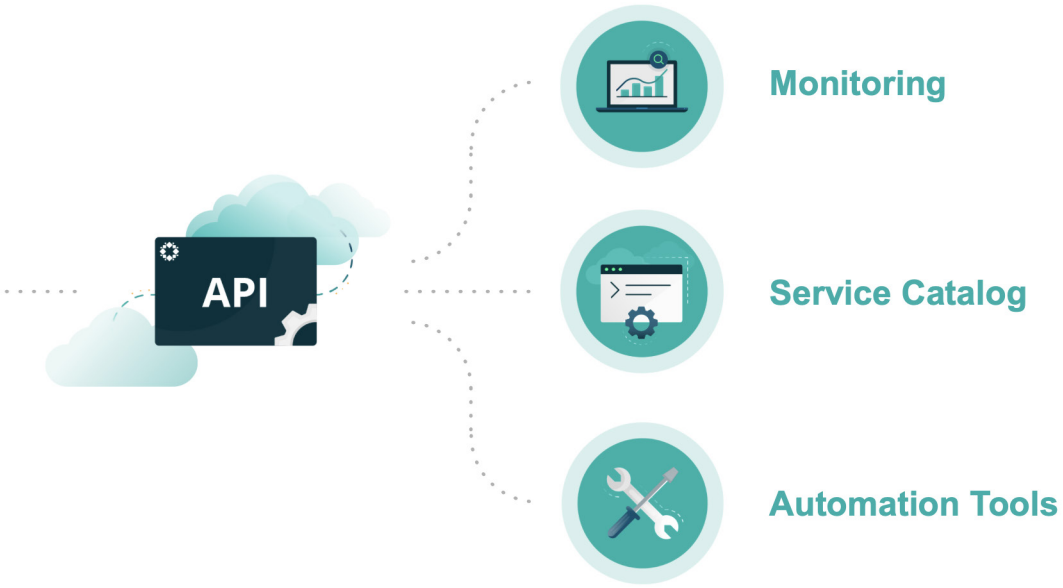

Maintain More Database Copies with Less Storage and Cost

And, you no longer have to waste precious storage capacity for each copy. By accessing backups directly, modern data protection can present users with "virtual copies" instantly. Users can gain access to up-to-date database copies whenever they are needed, for as long as needed, eliminating friction that slows business progress and freeing up DBAs for other value-added tasks.

Since no data is copied, a user can gain access to their own "copy" of a multi-terabyte database in just minutes, and storage is only consumed by metadata and user changes, so the incremental cost is near zero. Copies can also be:

- Refreshed regularly, so decisions and testing aren't based on outdated data sets
- Created to any point in time for troubleshooting or problem reproduction

Managing Compliance

Compliance is critical in every industry. Organizations need to be able to demonstrate they are in compliance with internal policies for data protection as well as regulatory requirements for data privacy and security. This includes the ability to demonstrate that backups occur on schedule and that data is recoverable.

Legacy

If you're using backup processes similar to those described earlier—with a variety of scripts and schedules deployed to different servers—there's no central repository that shows all your backups and their status, or that alerts you when something goes wrong. In traditional database environments, visibility into what's going on in the data protection domain is largely restricted to DBAs. Managers only know what DBAs tell them, or what they can find out from spreadsheets and sporadic reports.

Demonstrating recoverability is a heavyweight task that can require DBAs to periodically set aside significant blocks of time (along with adequate compute and storage capacity) to validate recovery processes.

Modern

Modern data protection creates a central repository for all data protection information. From a single interface, you can quickly see if backups ran successfully and if you're satisfying internal guidelines and meeting SLAs for data retention, security, and so on. Managers can see what's happening firsthand using centralized dashboards and reports to quickly identify database recovery points. Regular reporting helps satisfy specific compliance requirements.

Modern data protection makes verifying recoverability a much lighter weight activity. Often, it's as simple as quickly mounting an existing backup and verifying that it's readable—no additional storage required. With full APIs, important verification tasks can be automated to ensure they don't get overlooked.

Key Takeaways

To summarize, here are some of the ways that a modern data protection solution can simplify database protection and deliver more value, while giving precious time back to DBAs and backup admins.

Reduce Manual Work

- Automate discovery of new database instances
- Protect data and meet SLAs with policy-based protection
- Enable self sesrvice so users can satisfy requests themselves

Control Recoveries

- RBAC lets backup admins manage DB backups while DBAs manage recoveries
- Automate and simplify recovery processes
- Achieve near-zero RTO by minimizing data movement

Ensure Recoverability

- See your entire database protection environment from one interface
- Verify backups at time of creation
- Mount backups directly to validate recoverability

Access the Cloud

- Mobilize your data to the cloud
- Archive to the cloud for lower cost
- Ensure cloud-resident apps are protected

About Rubrik for Database

Rubrik simplifies database protection, enabling you to orchestrate mission-critical data across both on-premises and cloud locations while unifying backup, recovery, archival, replication, search, analytics, compliance, and management in a single converged software platform.

With Rubrik, you'll be able to take immediate advantage of advanced features, including:

Automate Discovery and Protection

Eliminate painful scripting and job scheduling to free up DBA time

Full Control Over Recoveries

Choose from a range of automated recovery options or continue to leverage your own scripts

Near-Zero Recovery Time Objectives

Recover databases in minutes regardless of size with Rubrik Live Mount and Instant Recovery

Self-Service Clones

Create instant self-service clones without added storage or impacts to production

To find out how Rubrik can help you enhance data protection for your database environment and increase the productivity of your entire team, visit rubrik.com


rubrik